Spectral Sharpening by Spherical Sampling

Graham D. Finlayson,[1] Javier Vazquez-Corral,[2,3,*] , Sabine Süsstrunk[4], and Maria Vanrell,[2,3]

[1]*School of Computing Sciences, University of East Anglia, Norwich, NR4 7TJ, United Kingdom*

[2]*Computer Vision Center, Edifici O, Campus UAB, 08193, Cerdanyola del Vallès, Barcelona, Catalonia, Spain*

[3]*Computer Science Department, Universitat Autònoma de Barcelona, Edifici Q, Campus UAB, 08193, Cerdanyola del Vallès, Barcelona, Catalonia, Spain*

[4] *School of Computing and Communication Sciences, École Polytechnique Fédéral de Lausanne, Lausanne, Switzerland, CH-1015*

[*]*Corresponding author: jvazquez@cvc.uab.cat*

There are many works in colour that assume illumination change can be modelled by multiplying sensor responses by individual scaling factors. The early research in this area are sometimes grouped under the heading von Kries adaptation: the scaling factors are applied to the cone responses. In more recent studies, both in psychophysics and in computational analysis, it has been proposed that scaling factors should be applied to linear combinations of the cones which have narrower support: they should be applied to the so-called "Sharp Sensors".

In this paper we generalise the computational approach to spectral sharpening in three important ways. First, we introduce *Spherical Sampling* as a tool that allows us to enumerate in a principled way all linear combinations of the cones. This allows us to, second, find the optimal sharp sensors which minimises a variety of error measures including CIE Delta E (previous work on spectral sharpening minimised RMS) and colour ratio stability. Lastly, we extend the spherical sampling paradigm to the multispectral case. Here the objective is to model the interaction of light and surface in terms of colour signal spectra. Spherical sampling is shown to improve on the state of the art.

1. Introduction

Colour values are generated through the interaction between the reflectance of an object, the illumination of the scene and the spectral sensitivities of the eye or camera. A simple colour formation model is written as:

$$p_k = \int_\omega \mathcal{R}_k(\lambda)E(\lambda)S(\lambda)d\lambda, \; k = \{R, G, B\} \tag{1}$$

Here $E(\lambda)$ and $S(\lambda)$ are, respectively, the spectral power distribution of the light and the spectral reflectance of a surface (in percentage). $\mathcal{R}_k(\lambda)$ denotes one of the three spectral sensitivities of the visual system (or camera) and p_k denotes the R, G, or B response. The integral is taken over the visible spectrum ω, which here and throughout this paper is assumed to be between 400 and 700 nanometres. From Equation (1) it is not immediately clear how the response vector $\underline{p} = [p_R, p_G, p_B]$ for a given object reflectance under one light relates to those under a second light. Remarkably, the relationship is statistically well modelled by a 3×3 linear transform. We can write:

$$\underline{p}_i^1 \approx M^{1,2}\underline{p}_i^2 \tag{2}$$

where the single superscript denotes dependence on illumination and the subscript depends on surface. Equation (2) teaches that surface i viewed under illuminant 1 is a 3×3 matrix $M^{1,2}$ transformation from the corresponding response vector under illuminant 2. The matrix $M^{1,2}$ only depends on the illuminant pair and not on the reflectance.

There are many studies explaining why equation (2) works. Assuming surface reflectance functions can be written as a linear combination of 3 basis functions, then illuminant change is exactly modelled by a 3×3 transform [1]. This idea was generalised by Marimont and Wandell [2] who considered non-linear models of reflectances but linear models of the responses themselves. It could, for example, be the case that reflectance spectra are 5 or 6 dimensional [3] but even using these surfaces, interacting with typical lights, the model (2) still works (in this case Marimont and Wandell showed how to derive an "effective reflectance" basis, i.e., that part of real reflectances that is important for image formation). More formally, Forsyth [4] precisely defines the conditions where (2) will hold exactly. Empirically, Equation (2) suffices to model illuminant change for the human visual system and all commercial cameras the authors are aware of.

Equation (2) is of central importance to investigations into colour constancy. If we wish to discount the *colour* cast due to the prevailing light from an image, it suffices to find the mapping (the 3×3 matrix transform) that removes this colour cast. For example, if an image is captured under bluish light, then all the recorded sensor responses are biased in the blue direction and, in particular, a white surface will itself be bluish. If we can find the map that

2

takes us from the blue light to a white reference illuminant, then applying this map will remove the colour cast. This idea of finding the appropriate map is elucidated in Forsyth's colour constancy theory [4]. In fact, almost all constancy algorithms can be viewed from this perspective. Of course, as simple as (2) is, there are 9 components in a 3×3 matrix and so colour constancy, i.e., calculating the 3×3 matrix, is a 9-dimensional problem. Apart from a theoretical algorithm presented by Forsyth, there are no practical 9-dimensional colour constancy algorithms (there are algorithms which solve for the transform matrix but in each case, the transform is parametrised by fewer than 9 numbers).

Early on, researchers [5, 6] adopted the even simpler von Kries, or Diagonal, model of illuminant change:

$$\underline{p}_i^1 \approx D^{1,2} \underline{p}_i^2 \tag{3}$$

where $D^{1,2}$ is a 3×3 diagonal matrix. Typically, the components of the diagonal matrix are chosen so that "white looks right": if $\underline{w}^1$ and $\underline{w}^2$, respectively, denote the sensor response to a perfect white diffuser for the two lights, then

$$D_{i,i}^{1,2} = w_i^1 / w_i^2 \tag{4}$$

Equation (3) implies that illumination change is a process which operates in each sensor response channel independently. This greatly simplifies colour constancy computation and we now only need to estimate the 3 parameters of the diagonal matrix transform. Remarkably, the diagonal model turns out to be rather good at accounting for illuminant change in many circumstances. Partly, this is explained by the fact that as the support of the sensor becomes small, then a diagonal matrix will work well [4]. Empirically, a diagonal matrix works for most cameras that have spectral sensitivities with support of 100 to 150 nanometres [5]. We (and many others) have found that equation (3) works well for many camera systems.

In contrast with typical camera systems, cone spectral sensitivities are quite broad and Von Kries adaptation, used to model illumination change, performs relatively poorly. In Figure 1 we plot (top) the Smith-Pokorny [7] cone spectral sensitivities, (middle) the XYZ colour matching functions [8] (which are a linear combination of the cones and are used for measuring and communicating colour) and (bottom) the RGB sensitivities of a SONY DXC-930 digital video camera [9]. A diagonal matrix works well for modelling illumination change for the camera sensitivities shown (they meet the 150nm support condition) but only partially for the cones [7] or colour matching functions [8].

In Spectral Sharpening we seek to find a linear combination of the cones (or, indeed, any other set of spectral sensitivities) with respect to which a diagonal model of illuminant change works best. Mathematically, the Spectral Sharpening idea is written as:
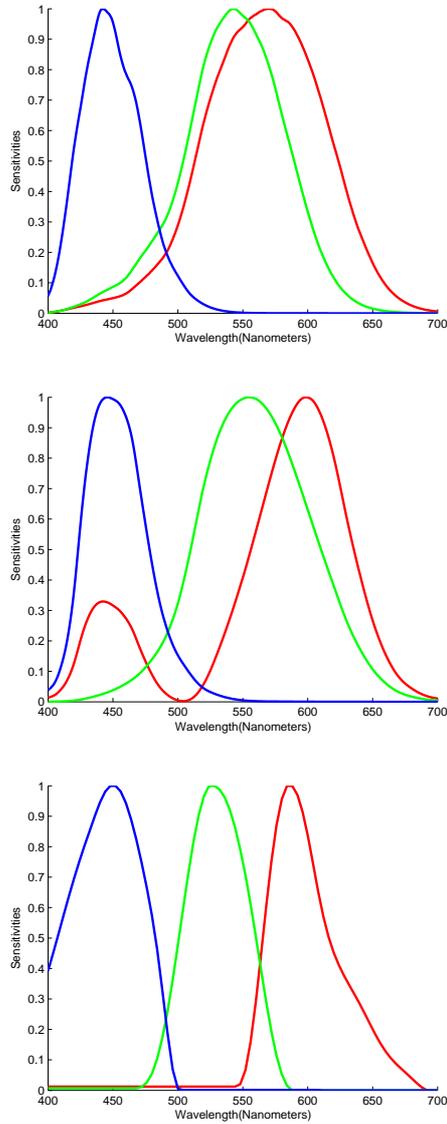
Fig. 1. Pokorny and Smith cones sensitivities (top) XYZ sensitivities (middle), SONY DXC-930 sensitivities (bottom)

$$\mathcal{T}\underline{p}_i^1 \approx D^{1,2}\mathcal{T}\underline{p}_i^2 \tag{5}$$

where the 3×3 matrix $\mathcal{T}$ is called a sharpening transform. The effect of linear combining cone responses is the same as substituting $\underline{\mathcal{R}}^\sharp(\lambda) = \mathcal{T}\underline{\mathcal{R}}(\lambda)$ for $\underline{\mathcal{R}}(\lambda)$ in the image formation equation (1). Rearranging (5):

$$\underline{p}_i^1 \approx \mathcal{T}^{-1}D^{1,2}\mathcal{T}\underline{p}_i^2 \tag{6}$$

The form of Equation (6) is useful for comparing the efficacy of diagonal models of illumination change: we begin in cone space, transform to sharp space, apply the von Kries scalings and finally transform back to cone space. Because we end up back in cone space, we are able to directly compare the performance of the diagonal model of illuminant change, either using cone sensors directly or sharp sensors. In both cases, cone responses under one light are mapped to cone responses under a second light.

In [10–13], it was shown there existed a linear combination of the cones with respect to which a diagonal mapping approximately accounted for illuminant change (matrices of the form $\mathcal{T}^{-1}D\mathcal{T}$ mapped responses across illumination with a small error).

However, almost without exception, the methods developed in previous papers minimised a *linear* error, then the efficacy of the transform was then evaluated according to a non-linear metric (e.g. percentage error or CIE Lab colour difference). We examine this disconnect in this paper. Surely, if we are interested in finding the sharp transform that minimises colour differences, we should explicitly solve the colour difference minimisation problem. More generally, in this paper, we are interested in solving for the best sharpening transform for any arbitrary non-linear function.

Let $f()$ denote an arbitrary non-linear mapping taking colour responses to perceptually relevant coordinates. Examples of $f()$ include the CIE Lab equations (see Appendix A) and the CIE Luv coordinates [8]. We are interested in finding the best sharpening transform $\mathcal{T}$ that minimises for all reflectances i and all lights j

$$f(\underline{p}_i^{ref}) \approx f(\mathcal{T}^{-1}D^{ref,j}\mathcal{T}\underline{p}_i^j) \tag{7}$$

In (7), the superscript ^{ref} denotes a reference lighting condition. In our experiments (discussed later in the paper) we use D65 [14] as our reference light. To find the best $\mathcal{T}$ it is useful to rewrite (7) as the explicit minimisation of an *error function*

$$\min_{\mathcal{T}} \quad err = \sum_j ||f(\underline{p}_i^{ref}) - f(\mathcal{T}^{-1}D_{\mathcal{T}}^{ref,j}\mathcal{T}\underline{p}_i^j)|| \tag{8}$$

where $||.||$ is an appropriate norm (for CIE Lab it is the usual L_2 norm). Other norms, such as CIEDE2000, can also be considered. The diagonal matrix $D_{\mathcal{T}}^{ref,j}$, which depends on the

transform $\mathcal{T}$, is the ratio of the response to a white surface under the lights ^{ref} and ^{j} (see (4)). By minimising (8) we find the best sensor basis with respect to which a von Kries transform best minimises CIE Lab error.

Von Kries adaptation - that illumination change can be modelled by a diagonal matrix - directly implies that the ratio of responses of an arbitrary pair of surfaces u and v should be illuminant invariant. Indeed, the invariance of colour ratios implies von Kries adaptation and vice versa [15]. In computer vision, colour ratios are used as light invariants where they subserve tasks such as object recognition [16] and image indexing [17]. Colour ratios are also at the heart of Land's retinex theory [6] of human colour vision.

To maximise colour ratio stability across illumination we could solve for:

$$\min_{\underline{t}_k} \quad err = \sum_j \left\| \frac{\underline{t}_k \cdot \underline{p}_u^j}{\underline{t}_k \cdot \underline{p}_v^j} - \frac{\underline{t}_k \cdot \underline{p}_u^{ref}}{\underline{t}_k \cdot \underline{p}_v^{ref}} \right\| \quad k \in \{r, g, b\} \tag{9}$$

Here, $\underline{t}_k$ denotes a single sharp sensor and '.' is the vector dot-product. That is, we look at ratio stability per sensor class: we separately optimise for the best red-, green- and blue sensors. However, implemented directly the calculated error for one sharp sensor is not easily interpretable (and so comparable) to those of a second sensor: the error measure does not have a natural scale. So, we modify our optimisation to minimise a percentage error. Let the vector $\underline{c}^j(\underline{t}_k)$ denote m ratios where ^{j} is the jth of N illuminants under which ratios are calculated and $\underline{t}_k$ denotes dependence on a particular sharp sensor. If there are M reflectances then there are $m = (M/2)(M-1)$ ratios in $\underline{c}$ (where given the pair of ratios a/b and b/a we always choose the ratio that is larger). The final modified ratio error function is written as:

$$\min_{\underline{t}_k} \frac{1}{N} \sum_{j=1}^N \frac{\|\underline{c}^{D65}(\underline{t}_k) - \underline{c}^j(\underline{t}_k)\|}{\|\underline{c}^{D65}(\underline{t}_k)\|} \tag{10}$$

The optimisations of (8) and (10) are not easy to solve: there are no *closed form* solution to these problems and none of the spectral sharpening methods hitherto minimised Delta E or ratio error.

In this paper, we develop a new optimisation technique, which we call *Spectral Sharpening by Spherical Sampling*, for finding the minima of arbitrary error functions. In our approach individual sensors are visualised as points on the surfaces of a 3-dimensional sphere and each set of sharp sensors corresponds to three points selected from the sphere's surface. By combinatorially sampling all triples of sensors and evaluating the error function for each triple, it is a simple matter (and it turns out to be feasible too) to enumerate all possible sets of sharp sensors. For each set we can compute (8) or (10) and then choose the sharp sensors that perform best overall.

Of course, this approach can only work, first, if the set of all sensors is sampled sufficiently finely and second, that the mapping that takes a point on the sphere and maps it to a linear

combination of cone sensors, is well-behaved (proximate points on the sphere should map to proximate sensors). Both of these technical issues are developed and addressed in section 3. In section 4, we show how spherical sampling can be used to find sharp sensors for minimising CIE Lab error (8) and ratio error (10) that significantly improves on the state of the art.

In [18], the 3-dimensional spectral sharpening idea was extended to model spectral interactions. There, light and reflectance spectra are represented using linear combinations of a small number of the same basis functions. Relative to this model, computing the multiplication of light and surface is expensive: for example, if we adopt a 6-dimensional basis then the multiplication of light and surface is a weighted sum of "6 choose 2"=15 product functions. The only exception occurs when the basis functions have disjoint spectral support. In this case there are only 6 product functions (and spectral computation is twice as fast). In [18], Spectral Sharpening was used to find functions that have the most disjoint supports. In section 5 of this paper, we show how the Spherical Sampling method can be used to find a disjoint spectral basis that delivers more accurate spectral computation than before. The paper concludes in section 6.

2.   Background: Linear Spectral Sharpening

It may seem surprising that a full 3×3 matrix model of illumination change (2) can be replaced by the 3 component diagonal matrix (5) albeit with respect to sharpened sensors. To understand why such an approximation is possible and why the approximation should also be tolerably accurate, it is useful to represent light and reflectances using linear basis models. Let us assume that surface reflectance functions and illumination spectra can be written as a linear sum of 3- and 2- dimensional basis functions, respectively.

$$S(\lambda) = \sum_{i=1}^{3} \sigma_i S_i(\lambda) \ , \ \ E(\lambda) = \sum_{i=1}^{2} \epsilon_i E_i(\lambda) \tag{11}$$

With respect to these conditions, it was shown in [11] that there existed a spectral sharpening transform where (5) and (6) held exactly: the error in (8) and (10) is exactly zero. Significantly, this approach, called *perfect sharpening*, is an algebraic result and holds for all sensor sets. 3- and 2-dimensional spectral models of Daylight illumination and natural reflectances provide a coarse, but tolerable, model of real lights and surfaces.

Linear model theory was extended by Marimont and Wandell [2] so that rather than modelling spectral quantities, they find linear models which best account for how spectra interact and then project to form RGB values. For daylight spectra, the $3 - 2$ model (3-d reflectance and 2-d light) provides a reasonable fit of cone response data. Further, in [12], it was shown that for the converse $2 - 3$ (2-d reflectance and 3-d light) model condition, there is a single unique spectral sharpening transform with respect to which von Kries adaptation is also a perfect vehicle for illuminant change.

More pragmatically, it seems natural to consider the sharpening problem given RGBs for the same surfaces viewed under different lights (rather than beginning with spectral models). Let A and B denote $3 \times N$ matrices of cone responses for N surfaces under two lights and M be the 3×3 least-squares matrix that minimises $||MA - B||$ (M is calculated using the Moore-Penrose inverse: $M = BA^t[AA^t]^{-1}$). Rather than M, we would like to have a matrix for the form shown in Equation (5). That is, we wish to break down M into a sharpening transform and a diagonal matrix relating sharpened RGBs.

Serendipitously, the standard eigenvector/eigenvalue decomposition of matrix M returns the sharpening transform $\mathcal{T}$ and diagonal matrix $\mathcal{D}$ we seek:

$$\mathcal{T}^{-1}D\mathcal{T} = M \qquad (12)$$

It was shown that when we minimise $||M'\mathcal{T}A - \mathcal{T}B||$, the least-squares solution is diagonal ($M' = D$). That is, this diagonal matrix transform is the *optimal* way to describe illuminant change when the eigenvectors $\mathcal{T}$ are used as the sharpening transform. *Data-based* sharpening returns the optimal *least-squares* sharp transform for any pair of lights.

More recently, Chong *et al.* [19] generalised data-based sharpening and presented a least-squares formulation where for a set of illuminants the sharpening transform that best models illumination change is found. However, one problem with the Chong *et al.* method is that it is a local minimiser and as such it may not actually find the best sharpening solution overall. This said, it is our experience that the Chong *et al.* method will almost always converge to the overall best least-squares solution. Significantly, we can make this judgement only because the Spherical sampling spectral sharpening method we develop in section 3 can find the global minimum. With the output of our method at hand, we verify the Chong *et al.* method has also converged to the correct answer.

A third way to find a sharpening transform is to consider only the properties of the sensor spectral sensitivities themselves. Indeed, we commented earlier that a diagonal model of illumination change for the SONY DXC930 camera sensors (bottom panel of Figure 1) worked well but that it works less well for the cones or XYZ colour matching functions. By visual inspection, the camera curves have much narrower support, and this is an important indicator of how well illumination change can be modelled using a diagonal matrix. Indeed, in the limit, sensors which are sensitive to single wavelengths of light (they are Dirac Delta functions) support perfect von Kries adaptation irrespective of any constraints placed on the dimensionality of surface reflectances.

In [11], *Sensor-based* sharpening was developed to find the linear combination of the cones that had maximal sensitivity in a given wavelength band. Optimising in short-, medium- and long-wavelength bands leads to sharp-blue, sharp-green and sharp-red sensors. Sensor-based sharp sensors, which were derived without statistical knowledge, support similar von

Kries performance to the sensors which optimise statistical functions. We will return to sensor-based sharpening in section 5.

Sharp transforms have also been derived from psychophysical experiments e.g. to predict corresponding colour data [20] or as part of a colour appearance model [21]. Later in the experimental section we will use the "Bradford" [20] sensors as representative of the sharp sensors found in psychophysical experiments.

3.   Non-linear Spectral Sharpening

There is no reason to assume that the "numbers" most important to perception are the raw sensor response values recorded by the cones. Indeed, we are not aware of any perceptual model of how we see that starts and ends with cone responses. Rather, these responses are transformed by linear and non-linear functions to arrive at values which correlate more closely with perceptual data. Thus, finding the best sharpening transform $\mathcal{T}$ so that the approximation in (6) works best (i.e. minimising the fit in terms of the cone responses) may not be the best approach if we are interested in modelling what is actually visually important to an observer.

Of course modelling what we see is far from being a solved problem. The best we can do is to try and find perceptual correlates relative to a given visual task. Indeed, much of the applied work in imaging builds on research on colour difference estimation. The aim there is to model how we see two similar colours and how we then judge the magnitude between colour pairs. The CIE Lab formulae [8], for example, apply a non-linear transform to L,M and S cone responses to form a corresponding L, a and b triplet. The Euclidean distance between pairs of *Lab* triplets more closely accounts for colour difference judgements made by observers. Significantly, CIE Lab is at the foundation of several methods in applied imaging that attempt to model how we gauge the perceptual distance between images [22, 23].

In this paper, we wish to find the best spectral transform that minimises CIE Lab error. The function $f()$ in (8) is the mapping from cone responses to Lab triplets.

CIE Lab is far from being the only colour space of interest. So, we are interested in developing a general solution to (8). In this paper we are also interested in computing the best sharp transform of calculating colour ratios. Ratios are at the heart of the Retinex theory of colour vision as well as, more pragmatically, being used in a variety of tasks in computer vision [16, 17]. In this second case, we are interested in the optimisation described in (10).

3.A.   Gradient Descent and Direct Search Methods

The starting point to solve either minimisation could be to use a sharpening transform derived from an antecedent sharpening algorithm (e.g. the Perfect sharpening or Data-based

sharpening described in section 2). These minimisations, while not tailored to the error functions in (8) or (10), could be used as an initial guess for future optimisations. Classically, we calculate derivatives of the error function with respect to the sharpening transform and then add a "delta" to the sharpening transform so that the error is *locally* reduced. By repeatedly perturbing the current best sharpening transform, the error is driven lower until a local minimum is found. In the case where the error function is not easy to differentiate, a direct-search method can be applied. The principle is similar. We search around the current guess to find a direction of change that reduces the error. However, the disadvantage of search methods is they tend to converge to the local minimum slower than when the derivative of the error function is known.

Unfortunately, gradient descent and direct search methods, by definition, are only guaranteed to find a local minima. While for some problems one can prove there is a single global minimum, in which case a search strategy will discover it, this is not true for our optimisations. Later we compare the spherical sampling approach to finding the best sharp sensors for maximising ratio stability to those discovered using a search-based optimisation. Our new method delivers sensors which deliver better ratio stability.

3.B. *Spherical Sampling*

We begin with the simple observation that the optimisation problem would be easy if we had a fixed number of sharpening transforms. If we have M transforms we can rewrite the minimisation of (8) as:

$$\min_x \quad err = \sum_j ||f(\underline{p}_i^{ref}) - f(\mathcal{T}_x^{-1} D_x^{ref,k} \mathcal{T}_x \underline{p}_i^j)|| \quad , \quad x \in 1, 2, \cdots, M \tag{13}$$

and (10) is rewritten as

$$\min_x \frac{1}{N} \sum_{j=1}^{N} \frac{||\underline{c}_x^{D65} - \underline{c}_x^j||}{||\underline{c}_x^{D65}||} \tag{14}$$

As before, $\underline{c}$ denotes all pairs of ratios for M surfaces. The superscript denotes dependence on illumination and the subscript x denotes the xth set of sharp sensors. This optimisation is run 3 times to find the best red-, green- and blue- sharp sensors that deliver the best ratio stability.

We now show how we *can* generate a reasonable set of all sharp sensors and, in so doing, validate these discrete optimisations. To develop our method, let us begin by thinking of our sensor spectral sensitivity functions as a 31×3 matrix $\mathcal{R}$ of numbers, i.e., we sample the L-, M- and S-cone mechanisms at 10 nanometer steps across the visible spectrum (400 to 700nm). In this discrete domain, a new sensor, which is a linear combination of the cones,

can be written as $\mathcal{R}\underline{p}$, where $\underline{p}$ is a 3-vector. Without loss of generality, let the magnitude of $||\underline{p}|| = 1$ (since the magnitude of a sensor has no bearing on the adequacy or otherwise of von Kries adaptation modelling illumination change). Accordingly, we can think of a new sensor as a point on the surface of a 3-dimensional sphere (see figure 7 for a picture of points uniformly distributed on a sphere). By randomly generating points, by spherical sampling, on the surface of the sphere we can, discretely, model all possible sensors, which are linear combinations of those represented by $\mathcal{R}$. Selecting a triplet of points from the surface of the sphere generates a set of 3 new sensors. Every triplet of points on the sphere's surface represents a candidate set of sharp sensors.

Ideally, we would like the distance between neighbouring points on the sphere to correspond to the distance between the generated sensors. If $\underline{p}$ and $\underline{p}'$ are coefficient vectors such that $||\underline{p} - \underline{p}'||^2 = \delta$, the squared distance between the corresponding sensors is written as:

$$[\underline{p} - \underline{p}']^t \mathcal{R}^t \mathcal{R}[\underline{p} - \underline{p}'] \neq \delta \tag{15}$$

To make the differences between coefficient vectors map to the same differences in sensors we need to use an orthonormal sensor basis rather than the cones. Let us decompose our 31×3 sensors $\mathcal{R}$ using the singular value decomposition (SVD) [24].

$$\mathcal{R} = \mathcal{U}\Sigma V^t \tag{16}$$

here $\mathcal{U}$ is the orthonormal matrix, Σ is a 3×3 diagonal matrix and V^t is 3×3 orthonormal matrix. Because $\mathcal{U}$ is orthonormal ($\mathcal{U}^t\mathcal{U} = \mathcal{I}$ [the identity matrix]), if $||\underline{p} - \underline{p}'||^2 = \delta$ then

$$[\underline{p} - \underline{p}']^t \mathcal{U}^t \mathcal{U}[\underline{p} - \underline{p}'] = \delta \tag{17}$$

that is, the difference between points on the sphere maps to the same difference in the generated sensors. Let us randomly choose 3 points on the sphere, $\underline{p}_1, \underline{p}_2, \underline{p}_3$. Then, we can define a new candidate sharp sensor set as:

$$\mathcal{R}^\sharp = \mathcal{U}P, \quad P = [\underline{p}_1, \underline{p}_2, \underline{p}_3]; \tag{18}$$

By randomly sampling all the points on the sphere, i.e., considering all possible P's of Equation (18), we effectively enumerate all possible sensor sets. The generated new sensor sets are related to the original cone basis $\mathcal{R}$:

$$\mathcal{R}^\sharp = \mathcal{U}P = \mathcal{U}\Sigma V^t(\Sigma V^t)^{-1}P = \mathcal{R}(\Sigma V^t)^{-1}P. \tag{19}$$

The sharpening matrix is written as:

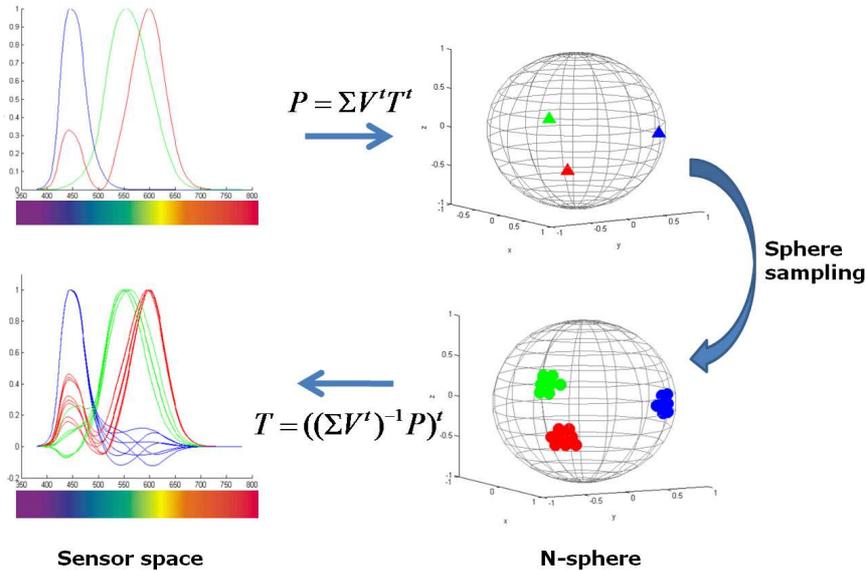$$\mathcal{T} = ((\Sigma V^t)^{-1}P)^t. \tag{20}$$

Fig. 2. Schema of the *spherical sampling* method. From a original sensor, we find the set of points P in the sphere. We sample the sphere around those points, and then return to the sensor domain.

Empirically, we have found that linear combinations of sensors to a resolution of about 1 degrees apart sufficed for optimising (13) or (14). This said, for the minimisation of (14) it is advantageous to use a finer search (computationally plausible since in that case we are treating each sensor separately). Figure 2 shows the relationship between sensors, points on the sphere and the sharpening transform. A 1-degree perturbation of XYZ functions is shown in Figure 3 where the original and new XYZ functions were scaled so that the maxima for each curve was 1.

Intuitively, a 1-degree resolution makes sense since if we cannot visually distinguish between two plots of a curve 1-degree apart, then it seems unlikely the difference is significant for vision.

3.C. *Determining the sampling resolution*

The sampling resolution we use is, however, not just based on intuition. Rather, we took the 1995 reflectances and 102 illuminants from [9] and generated *Lab* triplets for all $1995 * 102 \approx 200000$ cone response vectors. We then perturbed the true XYZs by k degrees and computed approximate *Lab* values for the perturbed XYZ colour matching functions. Comparing actual with approximate Labs we calculated the Delta E error to determine a measure of the visual significance of a k-degree sampling. We reduced k until the sampling error was in the range
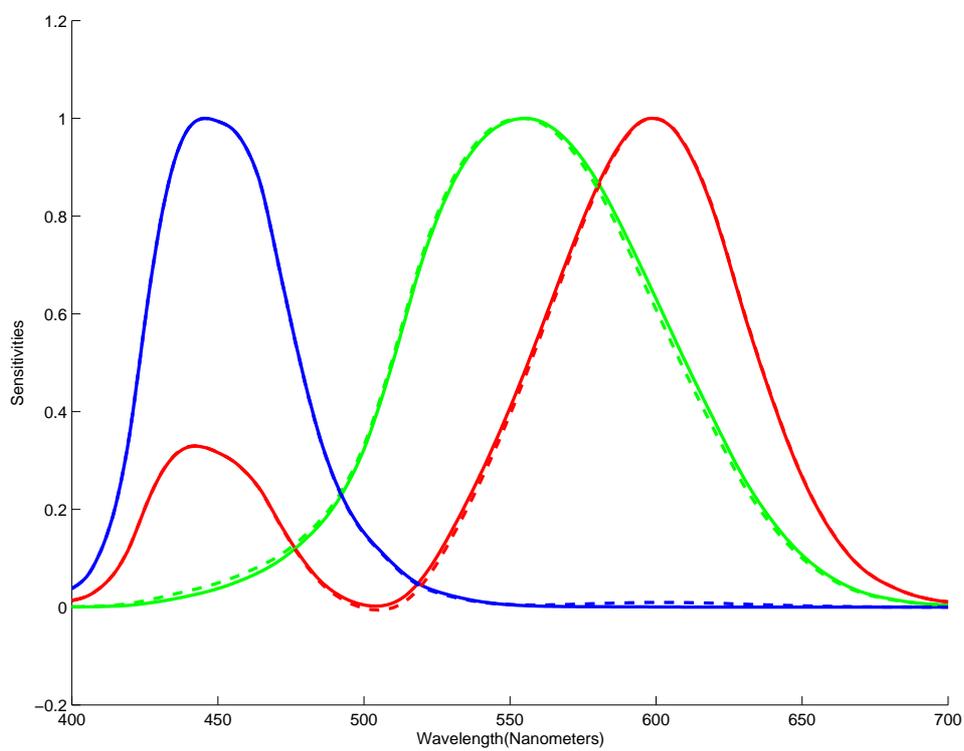
Fig. 3. XYZ curves (solid lines) versus XYZ perturbed 1-degree (dashed lines)

of about 1 Delta E unit. We found that for $k = 1$, the average was less than 0.8 and no single value was greater than 8. That is, images containing colours generated from actual or approximate XYZs would be indistinguishable to a human observer [25].

3.D. A cautionary remark

Suppose we repeat this experiment but now start with the cone fundamentals. We generate a second approximate set where each cone sensor is moved 1-degree away from the original. As before, we wish to consider the similarity of the sensors in terms of how they sample light and how similar or otherwise the responses are according to the CIE Lab formula. However, before we can use the CIE Lab formula, we must first map the cone responses to XYZs. Unfortunately, because long- and medium-wave sensitive cone sensors are similar in shape, the transform that maps cone curves to XYZ matching functions is in some mathematical sense *unstable* (specifically, the transform will have a high condition number). The import of this is that two cone response vectors for the same surface and light for the original and perturbed sets of cone sensitivities can be mapped to XYZs, which are much further apart. For this second experiment we would conclude that sets of cone sensors (one degree apart) were not "visually" equivalent to one another. A 1 degree sampling resolution would not suffice. Yet, so long as we are seeking sensors which are sufficiently decorrelated (true for the sharp sensors we seek) this instability problem cannot occur. However, care must always be taken to assess whether the problem can be placed in the discrete domain (e.g. whether for the optimisation at hand a 1 degree difference might be significant).

3.E. Implementation

We generate points on the surface of a 3-d sphere using Lovisolo and da Silva's method [26] (see Appendix B). In order to meet the expected 1 degree sampling resolution, we need to generate 25000 points on the sphere. Of course we are, potentially, interested in all triplets of sphere points: $\binom{25000}{3} = 2.6$ trillion sensor sets. If we treat $R(\lambda)$ and $-R(\lambda)$ as the same sensor then there are about 1.3 trillion sets of sensors. While this is a very large number, it is feasible to enumerate all these sensor sets.

However, for the spectral sharpening problem we do not have to consider all sensor sets. Indeed, while different studies have found a variety of sharp sensors, the discovered sensors were not so different from one another. So, we do not expect the best sharp set for any optimisation to be too far from any other set of sharp sensors. Thus, we propose using an answer from the linear optimisation approach (here we use Chong *et al.*'s algorithm [19]) as "seed" sharp sensors. We need only search for sensors that are nearby. Our sharp red, green and blue candidate sets are chosen so that all sensors are within 10 degrees of Chong's seed set, resulting in 185 sensors per channel or 6 million sensor sets in total.

4. Experiments

The spherical sampling approach can be used to implement all of the *linear* spectral sharpening methods discussed in section 2. Indeed, despite discretising the problem, spherical sampling (as expected) returns almost the same sensors as found by the analytic methods themselves. In this section we use spherical sampling to minimise the non-linear error functions of (13) and (14).

4.A.   Finding sharp sensors that minimise CIE Delta E

Here, we search for sensors that are optimal with respect to CIE Delta E. Let us calculate the colour response for a reflectance $S(\lambda)$ under a $D65$ illuminant [14] and the colour matching functions $\underline{\mathcal{X}}(\lambda)$. We calculate its CIE Lab representation $\underline{l}$:

$$\underline{l} = f(\underline{\rho}) = Lab(\underline{\rho}) \;\;,\;\; \underline{\rho} = \int_\omega S(\lambda)D65(\lambda)\underline{\mathcal{X}}(\lambda)d\lambda \tag{21}$$

Now, we compute the XYZ value for the same reflectance $S(\lambda)$ under a second arbitrary illuminant $E(\lambda)$:

$$\underline{\rho}' = \int_\omega S(\lambda)E(\lambda)\underline{\mathcal{X}}(\lambda)d\lambda \tag{22}$$

The XYZ response is mapped to the D65 reference illuminant using a sharp transformation:

$$\underline{\hat{\rho}} = \mathcal{T}^{-1}D^{D65,E}\mathcal{T}\underline{\rho}' \tag{23}$$

Here $\mathcal{T}$ is the matrix that maps XYZs (as opposed to cone fundamentals) to sharp counterparts. Also note that the matrix $D^{D65,E}$ has diagonal components equal to the sharp response of the sensors to D65 light divided by the sharp response to the second light E (see Equation (4)).

Now, let $\underline{\hat{l}} = Lab(\underline{\hat{\rho}})$ denote the estimate of the Lab coordinates of the surface under D65.

$$\Delta_E = \|\underline{l} - \underline{\hat{l}}\| \tag{24}$$

Now, we wish to use the minimisation of (13) to find the sharp transform that minimises Delta E errors. For reflectances, we used the 1995 reflectances of Barnard *et al.* [9]. We use four different illuminant datasets. The first dataset consists of the 102 illuminant set of SFU [9]. The other three datasets are subsets of this one. In dataset 2, there are 11 lights comprising three Fluorescent lights and 8 daylights. The third dataset contains just the daylights from dataset 2. Finally, dataset 4 was computed in the following way: We converted the 102 illuminants to CIELab. Afterwards, we computed a k-means clustering for 10 classes. Once the classes were constructed we selected the illuminant closest to the centroid of each class.

For each dataset we use our spherical sampling method to optimise for the best sharp sensors that minimises the mean, median and 75% quantile error. In all cases the Chong *et al.* sharpening transform is used as a seed (with respect to which we generate proximal sharp sensors by spherical sampling). Results are summarised in Table 1. The best algorithm for each column is shown in bold.

	Dataset 1			Dataset 2			Dataset 3			Dataset 4		
	Med	Mean	75%	Med	Mean	75%	Med	Mean	75%	Med	Mean	75%
Spherical S.	**0.79**	**1.79**	**2.00**	**0.93**	**2.23**	**2.50**	**0.58**	**1.07**	**1.20**	**0.91**	**2.17**	**2.59**
Chong's method	0.99	1.90	2.35	1.45	2.50	3.27	0.64	1.13	1.35	1.51	2.62	3.43
Bradford	0.95	1.99	2.26	1.21	2.59	2.97	0.83	1.34	1.71	1.07	2.42	2.95
XYZ	1.80	2.85	4.02	2.53	3.65	5.16	1.98	2.90	4.06	2.19	3.30	4.69

Table 1. Results of the experiment minimising CIE Delta E error for the four different illuminant datasets (measures: median, mean and 75% of Δ_E).

It is clear that *in all cases* spherical sampling delivers a significantly improved Delta E error. Compared with the Chong *et al.* method (designed to optimise linear error), the improvement can be as much as 50%.

In Figure 4 we show the derived optimal sharp sensors for each of the four illuminant datasets.

4.B.   Finding sharp sensors that maximise colour ratio stability

We use the same 4 data sets as before, but we are now interested in finding the sharp transform that maximises ratio stability. Using our sampling method, we can find the sensors that minimises (14). From a 250000 points sampling on the sphere (finer than before as stated in section 3.B), we chose the 4200 points which were within 15 degrees from a seed sensor. More sensors are used here because we found that the sharp sensors which worked best for optimising ratio stability were more different from the *linear* sharp sensors found in antecedent studies than was the case for the sensors that optimized CIE Lab error. Here we use the Bradford sensors for seed points since these sensors were found to have better ratio stability compared with Chong (though, we arrive at an almost identical result if we seed with Chong). In the first row of Table 2, we report the average of the ratio error for our method.

The second row shows the best performance for other sharpening methods (we chose the sharp sensors from Chong *et al.* and Bradford that returned the lowest average % error).
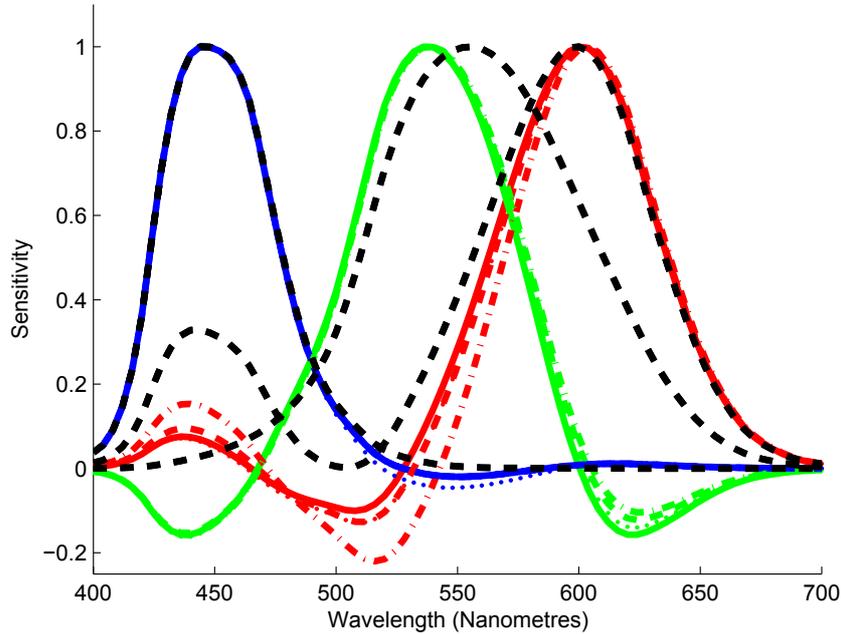
Fig. 4. Sharp sensors for the experiment minimising CIE Delta E error for the four different datasets. Dataset 1: 102 illuminants (straight line). Dataset 2: 8 daylights and 3 fluorescent lights (dashed line). Dataset 3: 8 daylights (dotted line). Dataset 4: 10 illuminants (dashed-dotted line). XYZ sensitivities are plotted in black

Again we note that the performance increase is quite large and is sometimes in excess of 50% compared to conventional sharpening methods.

In Figure 5 we show the sharp sensors that maximise ratio stability. Notice that compared to either the sensors found to minimise least-squares error or CIE Delta E, these sensors tend to be a little more broad-band. Intuitively this makes sense: the sharper the sensor the larger the negative lobe and the more likely that an integrated sensor response close to zero will be encountered. Small perturbations of small values can result in large ratio error. The sensors that support the most stable colour ratios have small negative lobes.

4.C. *Spherical Sampling vs Gradient Descent and direct search methods*

The reader may wonder whether a gradient descent or searching type approach [27] could be used so that similar results were achieved. To test this we took our seed sensors (the Bradford curves) and then simply searched for a local minimum of our error function. The results of this experiment are shown in the 3rd row of Table 2. The good news is that for
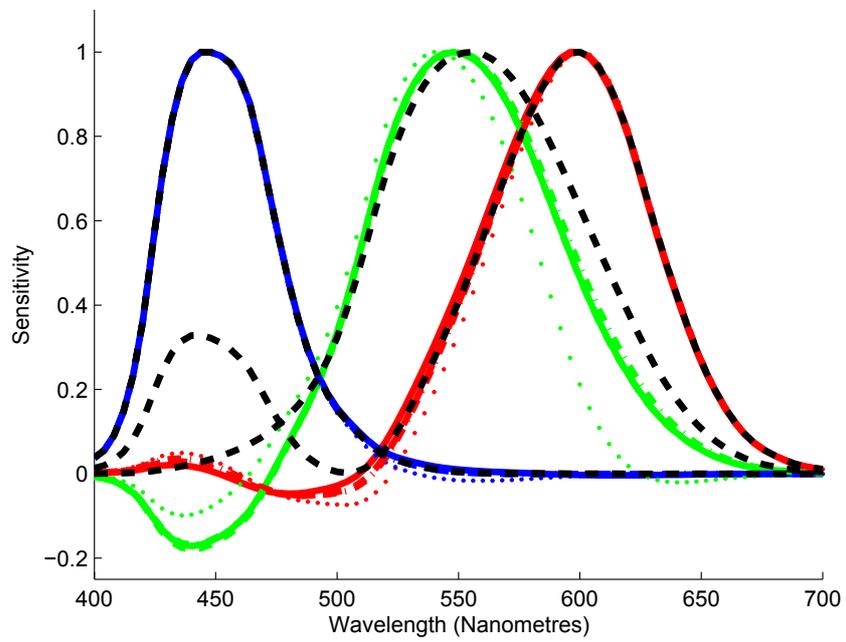
Fig. 5. Sharp sensors for the experiment maximising colour ratio stability for the four different datasets. XYZ sensitivities are plotted in black.

	Dataset 1			Dataset 2			Dataset 3			Dataset 4		
	R	G	B	R	G	B	R	G	B	R	G	B
Spherical S.	**0.0282**	**0.0739**	**0.0278**	**0.0354**	**0.0961**	**0.0326**	**0.0246**	**0.0360**	**0.0283**	**0.0322**	**0.0839**	**0.0317**
Other sharpening	0.0485	0.0861	0.0492	0.0685	0.1049	0.0404	0.0506	0.0393	0.0580	0.0707	0.1042	0.0381
Searching approach	**0.0282**	0.0747	0.0394	**0.0354**	0.1004	0.0727	**0.0246**	**0.0360**	0.0403	0.0337	*	0.0573

Table 2. Average ratio error for 4 illuminant data sets and 3 methods (see text for description)

4 of the 12 sensor optimisations, searching led to the same answer as spherical sampling (but no better). For 7 of the remaining search optimisations, the optimised sensor can be over 50% higher compared with spherical sampling. In the case where we optimise "sharp green" for the 4th illuminant dataset, the optimisation actually converges to "sharp red". This illustrates the usefulness of placing a prior constraint on the search space (as is easily done in spherical sampling).

5.   Finding sharp sensors that reduce the complexity of spectral calculations

Drew and Finlayson (henceforth DF) [18] showed that spectral sharpening could be used to simplify the cost of calculating the modelling of the interaction of light and reflectance spectra. Here we investigate whether spherical sampling can be used to further improve their method.

As in section 3, we carry out our development using 31-vectors to represent light and surfaces. Now, we write the product of light $\underline{E}$ of dimension (31×1) and surface $\underline{S}$, called the colour signal, denoted as $\underline{C}$

$$\underline{C} = diag(\underline{E})\underline{S} \qquad (25)$$

where $diag()$ converts a vector to a diagonal matrix.

DF define a matrix $\mathcal{C}$ $(31 \times K)$ containing the set of all possible colour signals given a set of lights and a set of reflectances. A low dimensional basis for this matrix $\mathcal{C}$ is computed via the singular value decomposition (SVD)

$$\mathcal{C} = U\Sigma V^t \qquad (26)$$

where U is a 31×31 matrix, Σ is a $31 \times K$ matrix and V is a $K \times K$ matrix. The first 6 columns of U defines a basis set B which is optimal in a least-squares sense. That is, linear combinations of the columns of B will best reconstruct the colour signal data set

(compared with any other choice of basis set). Suppose further that we describe reflectance and illumination using this basis (even though it was derived from their product). If $\underline{e}$ are the coordinates of an illuminant and $\underline{s}$ are the 6-dimensional coordinates for a reflectance in the basis B, then we would like $diag(\underline{e})\underline{s} = \underline{c}$ (where $\underline{c}$ is the coordinates of the actual colour signal spectrum in the same basis). DF call this the "factor model" of spectral image formation.

Intuitively, if B were a basis where the support of each function spanned a different part of the visible spectrum (at any wavelength only one basis vector is non zero) then this basis would behave in the desired manner. Yet, a basis defined in this way would not capture the statistics of lights and surfaces. Instead, DF proposed finding a linear combination of B where the factor model works best. They found their basis set using the "sensor-based" sharpening method (see [28] and background section).

Intuitively, if B contained 6 narrow band (power at a single wavelength) functions then the factor model would hold exactly, though, such a model would not accurately describe the continuous functions of wavelengths which are real colour signals. The idea in sensor based sharpening is to find the linear combination of the basis vectors that are maximally concentrated in a given wavelength band.

Let $\mathcal{T}$ denote the 6×6 matrix mapping the basis to a sharp counterpart:

$$\hat{B} = B\mathcal{T} \tag{27}$$

Let us denote the illuminant $\underline{E}$ with respect to the new basis by a 6×1 vector $\underline{e}$:

$$\underline{e} = \hat{B}^t \underline{E} \tag{28}$$

Similarly, $\underline{S}$ is represented by a 6×1 vector $\underline{s}$:

$$\underline{s} = \hat{B}^t \underline{S} \tag{29}$$

Given vectors $\underline{e}$, $\underline{s}$ the best least-squares approximation to the original signals is:

$$\hat{\underline{E}} = \hat{B}^+ \underline{e} \tag{30}$$

$$\hat{\underline{S}} = \hat{B}^+ \underline{s} \tag{31}$$

where $B^+ = B[B^tB]^{-1}$ (since $B^+B^t\underline{V} = B[B^tB]^{-1}B^t\underline{V}$ and $B[B^tB]^{-1}B^t$ is the *projector* of B).

In the new basis, a colour signal $\underline{c}$ is computed as:

$$\underline{c} = diag(\underline{e})\underline{s} \tag{32}$$

And, $\underline{C}$ is computed as:

$$\hat{\underline{C}} = \hat{B}^+ \underline{c} \tag{33}$$

5.A. Results

Let us repeat DF's spectral computation experiment. We calculate the best 6 colour signal basis vectors which model all pairs of 7 lights (A, C, D48, D55, D65, D75 and D100) multiplied by Vrhel's 170 measured reflectances [29]. We then carried out "sensor-based" sharpening to find the best "almost disjoint" linear combination that DF used in their factor model of spectral calculation. This factor basis is also used as the "seed" for further optimisation using our spherical sampling method. Figure 6 (top) shows the colour signal basis functions and Figure 6 (middle) the corresponding sharpened counterparts.

For a given approximate colour signal spectrum (calculated using the factor model) we could compare it to the true original in many ways. We could compare the spectral difference (one of the tests used by DF). Equally, and this is what we do here, we could integrate the spectra with the XYZ colour matching functions and then calculate the CIE Lab colour difference. We use 6-d spherical sampling to find the best factor model to minimise CIE Lab Delta E error.

Using spherical sampling we generate **just 12** points near each of the 6 DF basis vectors (using a 10 degrees distance). We do this because this implies $12^6 \approx 3000000$ new colour signal bases to test. Because we only have a few neighbours, the sampling resolution is coarser than for the 3-d spectral sharpening case. The optimal basis found by spherical sampling $\hat{B}$ is the one which minimises CIE Delta E and is shown in Figure 6 (bottom).

Results for the original DF method and DF+spherical sampling are reported in Table 3. Spherical sampling, even at a coarser resolution, in 6D delivers a significant increase in performance.

Method	All illuminants	Without A illuminant
Spherical sampling	**1.27**	**0.99**
Sensor-based sharpening	2.60	1.60
Colour Signal basis	30.70	30.23

Table 3. Results of the experiment for spectra recovering. Measure: mean Δ_E. First column: all the illuminants. Second columns: without illuminant A

While we have substantively improved the performance of the factor model of spectral image calculation, this experiment also illustrates the limits of spherical sharpening. In 6 dimensions we were only able to enumerate all "factor" bases in a very coarse way. It could well be that even better results would be possible if we could, combinatorially, sample the search space more finely.
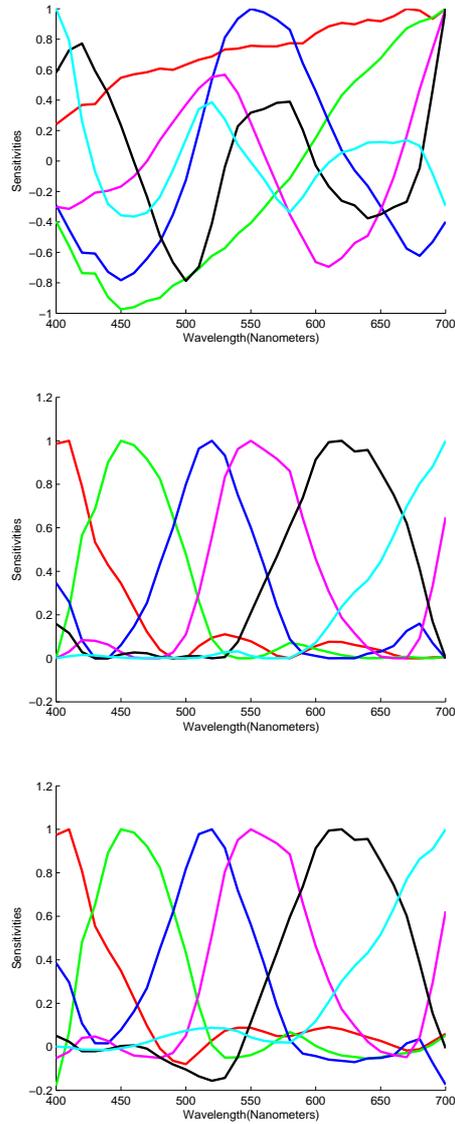
Fig. 6. Colour signal basis functions(top), DF corresponding sharpened counterparts (middle), Spherical sampling sharpened counterparts (bottom)

We believe that in 3-dimensional colour spaces spherical sampling will almost always suffice to find a good answer to a given optimisation. As the dimension of the optimisation problems increases the spherical sampling method will, like conventional search based methods, only be useful to probe "local" solutions.

6. Conclusions

Spectral sharpening finds linear combinations of the cones or camera sensors with respect to which a diagonal matrix accurately models illuminant change. There are many methods reported in the literature for finding sharp sensors. The difference in the methods is due to the difference in their optimisation criteria. However, in all cases, antecedent methods found the sharp sensors that minimised simple error measures (like linear least-squares).

In this paper, we present a new sharpening method that allows us to optimise an arbitrary objective function. We begin by showing how the set of all possible sensors for a trichromatic system can be represented as points on a 3-dimensional sphere. Finding all triplets of points allows us to enumerate all possible sensor sets. We used this spherical sampling procedure to find the sharp sensors which (1) minimised CIE Delta E error and (2) maximised the stability of colour ratios across illumination. The optimal sharp sensors found delivered significant performance increase compared to sharp sensors derived using conventional optimisation approaches.

We also extended our sharpening method to the multispectral case. Here, we wish to find a multispectral *sharp* basis with respect to which we can both model the spectral shape of lights and surfaces and model their interaction by the multiplication of their respective model weights. Effectively, we extend von Kries adaptation to the multispectral domain where it is called the *factor model of image formation*. Again, we show that spherical sampling can improve the basis selection. We find a basis set that leads to much improved perceptual spectral recovery using the factor model.

7. Acknowledgements

References

1. L. T. Maloney and B. A. Wandell, "Color constancy: a method for recovering surface spectral reflectance," Journal of the Optical Society of America A **3**, 29–33 (1986).

2. D. H. Marimont and B. A. Wandell, "Linear models of surface and illuminant spectra," Journal of the Optical Society of America A **9**, 1905–1913 (1992).

3. J. Park, M. Lee, M. D. Grossberg, and S. K. Nayar, "Multispectral imaging using multiplexed illumination," International Conference on Computer Vision (2007).

4. D. A. Forsyth, "A novel algorithm for color constancy," International Journal of Computer Vision **5**, 5–35 (1990).

5. J. A. Worthey and M. H. Brill, "Heuristic analysis of von Kries color constancy," Journal of the Optical Society of America A **3**, 1708–1712 (1986).

6. E. Land, "The retinex," American Scientist **52**, 247–264 (1964).

7. V. C. Smith and J. Pokorny, "Spectral sensitivity of the foveal cone photopigments between 400 and 500 nm," Vision Research **15**, 161 – 171 (1975).

8. G. Wyszecki and W. Stiles, *Color science: concepts and methods, quantitative data and formulae* (John Wiley & Sons, 1982), 2nd ed.

9. K. Barnard, L. Martin, B. Funt, and A. Coath, "A data set for colour research," Color Research and Application **27**, 147–151 (2002).

10. G. D. Finlayson, M. S. Drew, and B. V. Funt, "Color constancy: Enhancing von Kries adaptation via sensor transformation," in "Human Vision, Visual Processing, and Digital Display IV," (1993), pp. 473–484.

11. G. D. Finlayson, M. S. Drew, and B. V. Funt, "Spectral sharpening: sensor transformations for improved color constancy," Journal of the Optical Society of America A **11**, 1553–1563 (1994).

12. G. D. Finlayson, M. S. Drew, and B. V. Funt, "Color constancy: generalized diagonal transforms suffice," Journal of the Optical Society of America A **11**, 3011–3019 (1994).

13. G. Finlayson, *Coefficient Colour Constancy* (PhD thesis, Simon Fraser University, 1995).

14. J. Schanda, *CIE Colorimetry* (John Wiley & Sons, Inc.), pp. 25–78.

15. J. von Kries(1905), D. L. MacAdam, editor, Sources of Color Science pp. 110–119 (1970).

16. S. K. Nayar and R. M. Bolle, "Reflectance based object recognition," International Journal of Computer Vision **17**, 219–240 (1996). 10.1007/BF00128232.

17. B. V. Funt and G. D. Finlayson, "Color constant color indexing," IEEE Transactions on Pattern Analysis and Machine Intelligence **17**, 522–529 (1995).

18. M. S. Drew and G. D. Finlayson, "Multispectral processing without spectra," Journal of the Optical Society of America A **20**, 1181–1193 (2003).

19. H. Chong, S. Gortler, and T. Zickler, "The von Kries hypothesis and a basis for color

constancy," in "International Conference on Computer Vision," (2007).

20. K. Lam, *Metamerism and Colour Constancy* (PhD thesis, University of Bradford, 1985).

21. M. D. Fairchild, *Color Appearance Models* (Wiley-IST, 2005), 2nd ed.

22. Y. V. der Haeghen and J. M. Naeyaert, "Consistent cutaneous imaging with commercial digital cameras," Archives of Dermatology **142**, 42–46 (2006).

23. E. Chorro, E. Perales, D. de Fez, M. J. Luque, and F. M. Martínez-Verdú, "Application of the s-cielab color model to processed and calibrated images with a colorimetric dithering method," Optics Express **15**, 7810–7817 (2007).

24. G. H. Golub and C. F. V. Loan, *Matrix Computations* (The Johns Hopkins University Press, 1996), 3rd ed.

25. J. E. Gibson, M. D. Fairchild, and S. L. Wright, "Colorimetric tolerances of various digital image displays." in "Color Imaging Conference'00," (2000), pp. 295–300.

26. L. Lovisolo and E. da Silva, "Uniform distribution of points on a hyper-sphere with applications to vector bit-plane encoding," IEE Proceedings - Vision, Image, and Signal Processing **148**, 187–193 (2001).

27. J. C. Lagarias, J. Reeds, M. Wright, and P. Wright, "Convergence properties of the nelder–mead simplex method in low dimensions," SIAM Journal on Optimization **9**, 112–147 (1998).

28. M. S. Drew and G. D. Finlayson, "Spectral sharpening with positivity," Journal of the Optical Society of America A **17**, 1361–1370 (2000).

29. M. Vrhel, R. Gershon, and L. Iwan, "Measurement and analysis of object reflectance spectra," Color Research and Application **19**, 4–9 (1994).

Appendix A: Conversion from XYZ to CIELAB

Let us consider (X, Y, Z) a point under XYZ coordinates. Let us define the reference white point of XYZ as (X_w, Y_w, Z_w). The conversion from XYZ to CIELAB is computed as:

$$L^* = 116 f\left(\frac{Y}{Y_w}\right) - 16$$

$$a^* = 500\left(f\left(\frac{X}{X_w}\right) - f\left(\frac{Y}{Y_w}\right)\right) \tag{34}$$

$$b^* = 500\left(f\left(\frac{Y}{Y_w}\right) - f\left(\frac{Z}{Z_w}\right)\right)$$

$$\tag{35}$$

where

$$f(x) = \begin{cases} x^{\frac{1}{3}} & \text{if} \quad x > \frac{6}{29}^3 \\ \frac{1}{3}\left(\frac{29}{6}\right)^2 x + \frac{4}{29} & elsewhere \end{cases} \tag{36}$$

Appendix B: Lovisolo and da Silva method

The method proposed by Lovisolo and Da Silva [26] is based on the following assumption. If we have a big number of points K, there is one set of vectors uniformly distributed on a n-dimensional hyper-sphere that defines a tiling on the hyper-sphere by identical hypercubes of dimension $N - 1$.

A point p over the n-sphere in spherical coordinates will be represented as $p = [1 \ \omega_1 \cdots \omega_n]$. If we define the length of each hypercube as δ, this value can be defined as an addition of small variation of this spherical coordinates $\Delta_{\omega_1}, \cdots, \Delta_{\omega_n}$ as follows

$$\delta = \Delta_{\omega_1}$$

$$\delta = \Delta_{\omega_j} \sum_{i=1}^{j-1} \sin\omega_j, \ j = 2 : n - 1 \tag{37}$$

It should be noted that, given δ and a point $\omega_1, \cdots, \omega_n$ we can compute the different values for $\Delta_{\omega_1}, \cdots, \Delta_{\omega_n}$. On the other side, it is possible to compute the length δ assuming that the sum of the area of all the hypercubes shall be equivalent to the area of the sphere. From this, we have

$$A_n = K\delta^{n-1} \tag{38}$$

where K is the number of hypercubes or points we want to sample, and A_n is the area of the sphere, which can be computed as:

$$A_n = \frac{n\pi^{n/2}}{(n/2)!} \quad \text{for } n \text{ even}$$

$$A_n = \frac{n2^n\pi^{(n-1)/2}\left(\frac{n-1}{2}\right)!}{n!} \quad \text{for } n \text{ odd} \tag{39}$$

Using the relations stated above, Lovisolo and Da Silva proposed the following algorithm

1: compute Δ_{ω_1} from equation (37)
2: **for** $\omega_1 = \frac{\Delta_{\omega_1}}{2}$ to π in increments of Δ_{ω_1} **do**
3:    compute Δ_{ω_2} from equation (37)
4:    **for** $\omega_2 = \frac{\Delta_{\omega_2}}{2}$ to π in increments of Δ_{ω_2} **do**
5:     ...
6:     compute $\Delta_{\omega_{n-2}}$ from equation (37)
7:     **for** $\omega_{n-2} = \frac{\Delta_{\omega_{n-2}}}{2}$ to π in increments of $\Delta_{\omega_{n-2}}$ **do**
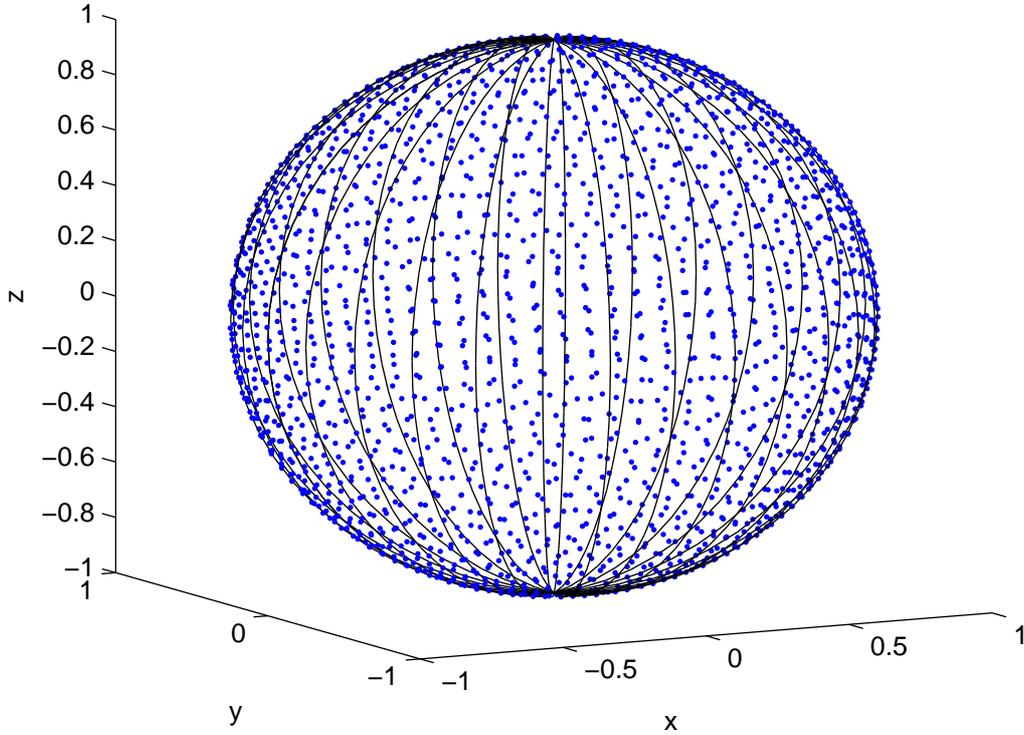
Fig. 7. Points sampled with the Lovisolo and Da Silva algorithm for three dimensions

8:           compute $\Delta_{\omega_{n-1}}$ from equation (37)

9:        **for** $\omega_{n-1} = \frac{\Delta_{\omega_{n-1}}}{2}$ to 2π in increments of $\Delta_{\omega_{n-1}}$ **do**

10:          **for** $i = 1...n - 1$ **do**

11:            compute $x_i = \sin \omega_{w_1} \cdots \sin \omega_{w_{i-1}} \cos \omega_{w_i}$

12:          **end for**

13:          compute $x_n = \sin \omega_{w_1} \cdots \sin \omega_{w_{n-2}} \sin \omega_{w_{n-1}}$

14:        **end for**

15:      **end for**

16:    **end for**

17: **end for**

An example of points sampled with the Lovisolo and Da Silva algorithm, 2000 points in the 3-d sphere, is plotted in Figure 7.